Module 16 – An Internet Exchange Point

Objective: To investigate methods for connecting to an Internet Exchange Point.

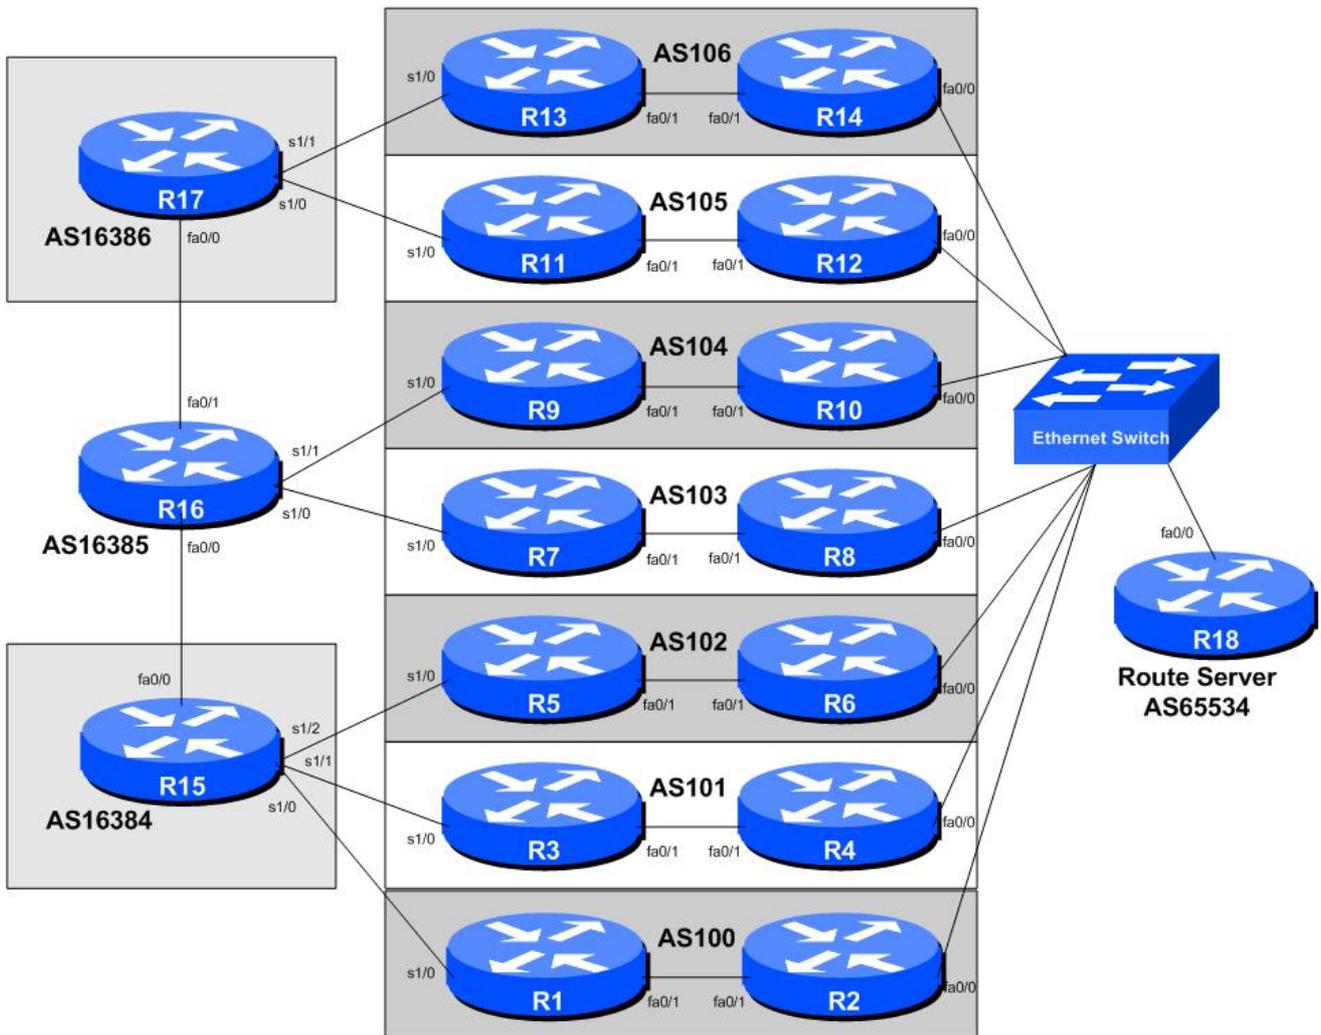**Prerequisites: Modules 12 and 13, and the Exchange Points Presentation**

Figure 1 – IXP Configuration

Lab Notes

The purpose of this module is to introduce the concept of an Internet Exchange Point, how to peer at IXPs, and look at some of the recommended configuration practices.

It has seven ASNs participating at the IX, with two routers being assigned to each autonomous system. One router peers at the IX, the other router is internal to the autonomous system (so will have an ISIS and iBGP session with the IX facing router). It also has three routers operating as three transit provider or Tier1 ISPs. This version of the module is incremental over Module 10 and covers correct ISIS and iBGP configuration practices for connecting the IXP facing router to the rest of the ISP network.

Lab Exercises

1. **A More Advanced IXP.** This IXP example has seven participating ASNs with each ASN having two routers. One router peers at the IXP; the other router is internal to the ISP's network and connects to an upstream ISP. The two routers in each ASN present at the IXP use ISIS and iBGP to communicate routing information with the each other. The three upstream ASNs are also interconnected with each other, to represent the high speed Internet core. What we have done here is to simulate an Internet Exchange Point and its participants' connections to upstream ISPs.

2. **Basic Configuration.** Each router team should configure their router to fit into the network layout depicted in Figure 1. Check all connections. Note that all links to the IX switch are by Ethernet. The links between each router in each AS is by serial connection (to represent the WAN connection from the IX back to the ISP's core network). The links from each AS to their upstream is by Ethernet (to represent the WAN link to the upstream).

3. **Addressing Plan.** These address ranges should be used throughout this section of the module. You are welcome to use your own range within an AS if you desire, just so long as you consult with the teams in other ASes to ensure there is no overlap.

AS100	**10.1.0.0/20**	**AS105**	**10.6.80.0/20**
AS101	**10.2.16.0/20**	**AS106**	**10.7.96.0/20**
AS102	**10.3.32.0/20**	**AS16384**	**10.8.112.0/20**
AS103	**10.4.48.0/20**	**AS16385**	**10.9.128.0/20**
AS104	**10.5.64.0/20**	**AS16386**	**10.10.144.0/20**

4. **Basic Router Setup.** Set up the routers as you would have done in previous modules. That is, basic security, IP addressing etc. The addressing for the point-to-point link between the two routers in each AS is left as an exercise to the Router Teams. But remember, a point-to-point link requires a /30 address block. Don't forget to configure the loopback interfaces too.

5. **IXP LAN.** The address range used for the IXP LAN is 172.17.10.0/24 – the route server (if provided in this module) has an IP address of 172.17.10.254. Each of the ASes is assigned addresses sequentially for use on the exchange point LAN. So, for example, AS100 has 172.17.10.1; AS101 has 172.17.10.2; AS102 has 172.17.10.3. And so on.

6. **Configure the ethernet of each router at the IXP.** The ethernet interfaces connected to the IXP should be configured appropriately for a public connection. Review the IOS Essentials materials and the IXP presentation. The configuration for Router 2 might be:

```
interface fastethernet 0/0
 description Exchange Point LAN
 ip address 172.17.10.1 255.255.255.0
 no ip directed-broadcast
 no ip proxy-arp
 no ip redirects
!
```

If you are unclear as to what any of the configuration lines do, please ask the lab instructor.

__Checkpoint #1:__ When you have properly configured your router, and the other routers at the IXP are reachable (i.e. you can ping the other routers), please let the instructor know.

7. **Configuring ISIS (part 1).** Obviously each IXP team will need to configure ISIS between the two routers in their AS. The router which is internal to the ASN's network should be straightforward to configure. It has one loopback interface, one interface connecting to the router at the IXP, and one interface connecting to the upstream. ISIS should only be activated on the internal interface, with the loopback marked as passive. Note we do not configure the interface pointing to the upstream as we will be using *next-hop-self* for our iBGP sessions. Example for Router 11 might be:

```
router isis as105
 net 49.0001.0100.0609.6001.00
 is-type level-2-only
 metric-style wide level-2
 set-overload-bit on-startup wait-for-bgp
 log-adjacency-changes all
 metric 100000
 passive-interface loopback0
!
interface fastethernet 0/1
 description BackBone link to R12
 ip router isis as105
 isis metric 2 level-2
!
```

8. **Configuring ISIS (part 2).** The ISIS configuration on the router connecting to the IXP needs more care. It is highly important that the IXP LAN network **is not** advertised into the ISP's iBGP. If it does, there is the possibility of leaking the IX LAN network by the ISP's BGP to other ASNs, providing transit to the IXP. This might not be a good thing, and many IXPs have rules against this behaviour. What we do want to do is include the IX LAN in the ISIS, and like the loopback interface, we mark the IX LAN interface as passive in ISIS.

The ISIS configuration for the IXP router looks something like this, for example for Router 8:

```
router isis as103
 net 49.0001.0100.0404.8001.00
 is-type level-2-only
 metric-style wide level-2
 set-overload-bit on-startup wait-for-bgp
 log-adjacency-changes all
 metric 100000
 passive-interface Loopback0
 passive-interface fastethernet0/0
!
```

```
interface fastethernet 0/1
 description Backbone link to R7
 ip router isis as103
 isis metric 20 level-2
!
```

NB. There is **NO** ISIS configuration applied to the IXP facing Ethernet interface. Do **NOT** put any in, repeat, do **NOT**! It is **NOT** a mistake.

9. **Configuring iBGP between upstream facing routers and the IXP facing router.** Any router which is connected to other ASNs (the ISP border routers) such as the ISP's upstreams or ISP's peers needs to have a modified iBGP configuration when peering with the IXP facing router. If all the prefixes learned from the upstream ISP are passed onto the IXP facing router, then there is the possibility that the IXP peers (and non-peers) can point static routes for those destinations at the IXP facing router, thereby gaining outbound transit across the local ASN. This is undesirable, and is basically a security risk. The configuration in the previous step therefore needs to be modified for IXP facing routers (the only router in this lab module example) so that external prefixes are not announced to the IXP facing router. In real life, ISPs create a peer-group which is used only when peering with IXP-facing routers, for example:

```
router bgp 103
 no synchronization
 bgp deterministic-med
 network 10.4.48.0 mask 255.255.240.0
 neighbor ibgp-ixp-br peer-group
 neighbor ibgp-ixp-br remote-as 103
 neighbor ibgp-ixp-br update-source loopback0
 neighbor ibgp-ixp-br next-hop-self
 neighbor ibgp-ixp-br password cisco
 neighbor ibgp-ixp-br description Used for iBGP with IXP BRs
 neighbor ibgp-ixp-br send-community
 neighbor ibgp-ixp-br filter-list 10 out
 neighbor router8-loopback peer-group ibgp-ixp-br
 distance bgp 200 200 200
 no auto-summary
!
ip route 10.4.48.0 255.255.240.0 null0
!
ip as-path access-list 10 permit ^$
!
```

Note the addition of the as-path access-list 10 – this basically allows only locally originated prefixes to reach the IX facing router. For example, Router7 with this configuration will now no longer send externally learned prefixes to Router8, thereby guaranteeing the security of the AS103 network.

10. **Configuring iBGP on IXP facing routers (Part 1).** The iBGP configuration on the IXP facing router needs much more care. We do **NOT** want the IXP LAN to appear in our iBGP. We again use the `next-hop-self` BGP configuration as we have done for our other iBGP sessions so that the IXP LAN is not required as an external nexthop – as we did earlier on. For example, for Router 6:

```
router bgp 102
 no synchronization
 bgp deterministic-med
```

```
    neighbor ibgp-peers peer-group
    neighbor ibgp-peers remote-as 102
    neighbor ibgp-peers update-source loopback0
    neighbor ibgp-peers next-hop-self        ! set external next-hops to us
    neighbor ibgp-peers password cisco
    neighbor ibgp-peers description Standard iBGP peers
    neighbor ibgp-peers send-community
    neighbor router5-loopback peer-group ibgp-peers
    distance bgp 200 200 200
    no auto-summary
    !
```

This configuration is also valid for other iBGP peers in the network if there were more than two routers in the ASN. ISPs try and keep configuration differences to the absolute minimum, and having two different peer-groups defined for iBGP is usually sufficient!

11. **Configuring iBGP on IXP facing routers (Part 2).** The second important difference required for the iBGP at the IXP router is that the local address block must **NOT** be originated here. If the IX router is disconnected from the core backbone for whatever reason, if it originated the ISP's address block, it would end up blackholing all traffic from the IXP participants to the local network. This is highly undesirable, so the correct procedure is to originate the ISP address block only in the core of the network, not the edge, so that normal BGP failover rules can apply.

    Q. What do you think the normal BGP failover rules might be in this case?

    A. If the IXP facing router is disconnected from the core network, the iBGP fails and the ISIS fails. So the router no longer hears the announcement of the ISP's address block from the core, so no longer announces it to the IXP peers, its eBGP neighbours. When this happen, the other ASNs will use fall back paths (alternative to the IXP) to reach the local ASN. These are the normal BGP failover rules.

12. **Configuring iBGP on the IXP facing routers (Part 3).** The second very important requirement at an Internet exchange point is that a default route (or the routes not originated by the local ASN) should not be made available at the peering LAN. If this requirement is not followed, then it is possible for the Exchange Point participants to use the local network to get transit to the rest of the Internet.

    In this Module, each ASN will be receiving a default route from their upstream provider (a common scenario in the Internet today). This default route must **NOT** be made available at the ISP's Exchange Point router. There are two ways to do this. Either block the announcement of the default route by iBGP at the borders router which connect to the upstreams, or set up a static default route to the Null interface at the IXP router (or indeed do both!). For example:

    ```
    ip route 0.0.0.0 0.0.0.0 null0
    ```

    Then if any of the IXP participants point a default route to the local network, the traffic will simply be dumped in the Null interface of their IXP router. Only traffic for specific destinations which available in the routing table on the IXP router will be forwarded to the rest of the network. This is a very important **network security** requirement.

13. **Configuring eBGP on the IXP facing routers.** Next, eBGP needs to be set up on the IXP routers. Create a peer-group and apply that peer-group to each eBGP neighbour. A sample configuration for Router8 might be:

```
ip prefix-list myprefixes permit 10.4.48.0/20
ip prefix-list peer100 permit 10.1.0.0/20
..
ip prefix-list peer106 permit 10.7.96.0/20
!
router bgp 103
 no synchronization
 bgp log-neighbor-changes
 neighbor ixp-peers peer-group
 neighbor ixp-peers remove-private-AS
 neighbor ixp-peers prefix-list myprefixes out
 neighbor ixp-peers route-map set-local-pref in
 neighbor <router2> remote-as 100
 neighbor <router2> description Peering with AS100
 neighbor <router2> peer-group ixp-peers
 neighbor <router2> prefix-list peer100 in
..
 neighbor <router14> remote-as 106
 neighbor <router14> description Peering with AS106
 neighbor <router14> peer-group ixp-peers
 neighbor <router14> prefix-list peer106 in
 no auto-summary
!
route-map set-local-pref permit 10
 set local-preference 150
!
```

The configurations for the other routers will be similar to this one. All router teams will have done sufficient BGP configuration throughout this workshop to extrapolate from the above examples. If in any doubt, ask the lab demonstrator for assistance.

Note the prefix-lists. There is a per-peer inbound prefix-list. Some service providers only filter ASes – that has inherent dangers, and does not prevent against inbound leaking of prefixes incorrectly originated by the peer AS. But only filtering on prefixes doesn't scale, especially in larger IXPs with large participating service providers as they are frequently adding to the prefixes they announce. The Internet Routing Registry is usually used to solve this problem.

14. **Set up passwords on the eBGP sessions at the IXP.** Negotiate with each ASN a password which you can use on your BGP session with them. And then agree to cut the eBGP session over to using passwords such that the eBGP session does not fall over due to password mismatches (as in Module 2). An excerpt from Router 8's configuration might be:

```
router bgp 103
 ...
 neighbor <router10> password peer104
 ...
!
```

15. **Configuring eBGP between the routers in AS16384, AS16385 and AS16386.** The instructors will operate the three transit/Tier1 ASNs; they will now configure eBGP between each other. No filtering is expected or required between the three ASNs, so the eBGP configuration is quite straightforward, for example as on Router 15:

```
router bgp 16384
 no synchronization
 bgp log-neighbor-changes
 neighbor <router16> remote-as 16385
 neighbor <router16> password as16385-password
 neighbor <router16> description Peering with AS16385
 no auto-summary
!
```

16. **Configuring eBGP with the routers in AS16384, AS16385 and AS16386.** All IXP members must now configure eBGP sessions with the routers in the three transit ASNs. They must expect to receive the default route from their upstreams, and only send their address block to the upstream. An example configuration for Router 3 might be:

```
ip prefix-list myprefixes permit 10.2.16.0/20
!
router bgp 101
 no synchronization
 bgp log-neighbor-changes
 neighbor <router15> remote-as 16384
 neighbor <router15> description Peering with AS16384
 neighbor <router15> password as16384-password
 neighbor <router15> prefix-list default in
 neighbor <router15> prefix-list myprefixes out
 no auto-summary
!
```

17. **Connectivity Test.** Check connectivity throughout the IXP network. Each router team should be able to see all the other routers at the IXP. When you are satisfied that BGP is working correctly, try running traceroutes to check the paths being followed.

```
Router9#trace 10.7.112.1

  1 10.5.64.6 4 msec 4 msec 4 msec
  2 172.17.10.4 4 msec 4 msec 4 msec
  3 10.7.112.1 [AS 106] 4 msec *  4 msec
Router9#
```

__Checkpoint #2:__ *Once the BGP configuration has been completed, check the routing table and ensure that you have complete reachability over the entire network. If there are any problems, work with the other router teams to resolve those.*

STOP AND WAIT HERE

18. **Implementing a Route Server at the IXP.** With route server functionality added to Cisco IOS as from the 15.2 release, it is now possible to operate Router 18 as a route server (IXPs tend to use Linux based routing software called BIRD or Quagga). The lab instructors will have configured Router 18 to be a route server.

The following information is included for the information of each router team member only. For those in the workshop interested in setting up an Internet Exchange Point following the completion of the workshop, this section contains the information necessary to configure and operate a Cisco IOS based Route Server.

The Route Server has eBGP sessions configured with each member of the IXP in this workshop lab (in real life, the Route Server has an eBGP session configured with those members of the IX who wish to use its facilities.) Each eBGP session is configured with the "route-server-client" command included in the eBGP peer-group.

This is an excerpt from the configuration of Router 18 after it has been configured to be a route server:

```
router bgp 65534
 bgp log-neighbor-changes
 neighbor ixp-peers peer-group
 neighbor ixp-peers description Route Server peers
 neighbor ixp-peers password cisco
 neighbor 172.17.10.1 remote-as 100
 neighbor 172.17.10.1 peer-group ixp-peers
 neighbor 172.17.10.1 description AS100 peer
 neighbor 172.17.10.2 remote-as 101
 neighbor 172.17.10.2 peer-group ixp-peers
 neighbor 172.17.10.2 description AS101 peer
 neighbor 172.17.10.3 remote-as 102
 neighbor 172.17.10.3 peer-group ixp-peers
 neighbor 172.17.10.3 description AS102 peer
 neighbor 172.17.10.4 remote-as 103
 neighbor 172.17.10.4 peer-group ixp-peers
 neighbor 172.17.10.4 description AS103 peer
 neighbor 172.17.10.5 remote-as 104
 neighbor 172.17.10.5 peer-group ixp-peers
 neighbor 172.17.10.5 description AS104 peer
 neighbor 172.17.10.6 remote-as 105
 neighbor 172.17.10.6 peer-group ixp-peers
 neighbor 172.17.10.6 description AS105 peer
 neighbor 172.17.10.7 remote-as 106
 neighbor 172.17.10.7 peer-group ixp-peers
 neighbor 172.17.10.7 description AS106 peer
 !
 address-family ipv4
  neighbor ixp-peers route-server-client
  neighbor 172.17.10.1 activate
  neighbor 172.17.10.2 activate
  neighbor 172.17.10.3 activate
  neighbor 172.17.10.4 activate
  neighbor 172.17.10.5 activate
  neighbor 172.17.10.6 activate
  neighbor 172.17.10.7 activate
  distance bgp 200 200 200
 exit-address-family
!
```

Notice that in this simple route-server example, there are no BGP filters at all – it is left to the IXP member to do necessary filtering.

This is just a simple example. Most IXPs will insist on the route server doing the basic bogon filtering which has been discussed elsewhere during this workshop. Some IXPs operate their route servers to filter according to the information contained in the Internet Routing Registry.

19. **Set up eBGP session with Router 18.** Each team member should now configure their router to peer eBGP with the route server. We will leave the existing bi-lateral peering sessions in place for the time being. Here is an example configuration:

```
router bgp 103
...
 neighbor 172.17.10.254 remote-as 65534
 neighbor 172.17.10.254 description eBGP with the IX Route Server
 neighbor 172.17.10.254 password cisco
 neighbor 172.17.10.254 remove-private-AS
 neighbor 172.17.10.254 prefix-list myprefixes out
...
```

Notice that the route server is running in a private AS – there isn't really any need for it to use a public AS as the Server does not need to be directly visible outside of the IXP.

Note also that there is no inbound prefix filter. This is not normally the case for an IXP. Even if an ISP peers with the Route Server, they should still filter inbound BGP prefixes they receive via the Route Server, just in case a peer who uses the Route Server advertises prefixes they should not be advertising.

20. **Modifications to the ISP**'s **IXP routers to talk to the Route Server.** Route servers, as a rule, do not include their AS number in BGP announcements to their clients; they are simply there to distribute the BGP table they have learned from their clients. Historically RSd and routing software platforms such as BIRD and Quagga have always behaved this way when running in route server mode. Some IXPs used routers instead of these software routing implementations, and routers have not had the functionality to drop the local ASN, meaning that the Route Server's AS has appeared first in the AS path.

Cisco routers participating at IXPs with the software based route servers have long had a capability to accept peerings with a full Route Server AS – but it is not IOS's default behaviour. IOS's default behaviour is to check that all prefix announcements from an eBGP peer have the peer AS as the first AS in the AS path. To support an eBGP peering where the Route Server AS is not included in the AS path, the router needs the following configuration:

```
router bgp 103
 no bgp enforce-first-as
...
```

which stops the router rejecting the prefix updates if the peer AS is not first in the AS path. You will notice that before entering this command, the BGP peering will be up, but no prefixes will be received by the ISP's IXP router – and the router logs will have messages about malformed attributes being sent by the Route Server.

Checkpoint #3: *Now look at your BGP routing table. You should see two paths to each destination, one via the RS, the other via the bi-lateral BGP peer.*

21. **Removing non-RS BGP sessions.** Once the peering session has been established with the Route Server and prefixes from all ASNs are being received from the Route Server, then each router team can remove their direct eBGP configurations with each AS at the IXP. While many ISPs choose to continue with bi-lateral peering configuration as well as the Route Server peering, this step here shows what happens when just the RS is used for the eBGP sessions at the IX:

```
router bgp 100
 no neighbor 172.17.10.2
 no neighbor 172.17.10.3
 no neighbor 172.17.10.4
 no neighbor 172.17.10.5
 no neighbor 172.17.10.6
 no neighbor 172.17.10.7
!
```

22. **Local Preference on RS supplied prefixes.** The prefixes heard from the Route Server all have the default local preference of 100. What each team should do now is use the previous route-map set-local-pref an applied it to the BGP peering with the RS. For example:

```
router bgp 102
 neighbor 172.17.10.254 route-map set-local-pref in
```

and then refresh the BGP session with the route-server. All the prefixes learned from the route-server will now have the local-preference set to 150.

23. **Filtering RS supplied prefixes.** It is generally quite a good idea not to trust what the route-server is sending to you. Each team should now take the individual prefix filters they had for each neighbour and **replace** them with a combined prefix filter for the route-server. This way each AS is still filtering what it receives from each AS. Note that in practice, many IXPs operating a route-server will filter the prefixes they receive from each participant in the route-server, and several more will filter what they send to each participant as well. This all helps with the integrity of the IXP peering mesh. For example:

```
router bgp 104
 neighbor 172.17.10.254 prefix-list RS-prefixes in
!
ip prefix-list RS-prefixes permit 10.1.0.0/20
ip prefix-list RS-prefixes permit 10.2.16.0/20
ip prefix-list RS-prefixes permit 10.3.32.0/20
ip prefix-list RS-prefixes permit 10.4.48.0/20
ip prefix-list RS-prefixes permit 10.5.64.0/20
ip prefix-list RS-prefixes permit 10.6.80.0/20
ip prefix-list RS-prefixes permit 10.7.96.0/20
!
```

Don't forget to tidy up unused configuration by removing the individual prefix-lists used in the previous section for the individual (bi-lateral) peerings.

Checkpoint #4: *Now look at your BGP routing table. You should only see prefixes from each AS, but learned via the route server.*

24. Summary. This module has given example configurations as used by Internet Service Providers at Internet Exchange Points. It has concentrated on using prefix-lists only – more sophisticated configurations are possible by using as-path filters and BGP communities. These examples are left to the reader to consider. If there is time at the end of the workshop, ask the Instructor to test out some other scenarios.